Fast extraction of the backbone of projected bipartite networks to aid community detection

J. Liebig[a], A. Rao[a]

[a]*School of Mathematical and Geospatial Sciences, RMIT University, Melbourne 3001, Australia*

arXiv:1512.01883v2 [cs.SI] 21 Mar 2016

Abstract

This paper introduces a computationally inexpensive method of extracting the backbone of one-mode networks projected from bipartite networks. We show that the edge weights in the one-mode projections are distributed according to a Poisson binomial distribution and that finding the expected weight distribution of a one-mode network projected from a random bipartite network only requires knowledge of the bipartite degree distributions. Being able to extract the backbone of a projection is highly beneficial in filtering out redundant information in large complex networks and narrowing down the information in the one-mode projection to the most relevant. We demonstrate that the backbone of a one-mode projection aids in the detection of communities.

Keywords: Bipartite network, One-mode projection, Backbone, Degree distribution, Community structure

1. Introduction

Scientist are often faced with an overwhelmingly large amount of data. Being able to reduce this to the most relevant information greatly simplifies analyses. In network science, a way of reducing networked data is discarding redundant or insignificant edges, resulting in a network that is commonly referred to as the backbone. The backbone of a network $\mathcal{G}(V, E)$, with node set V and edge set E, is defined as the subgraph $\mathcal{G}'(V, E')$ of $\mathcal{G}$, such that the edge set E' of the backbone $\mathcal{G}'$ contains only the most significant edges in E. Despite this straightforward definition, many challenges lie in the identification of significant edges.

Many different approaches to determining the most significant edges of a given network have been proposed, all pursuing the same goal of retaining only the most relevant information [1, 2, 3, 4]. Most methods focus on weighted one-mode networks, identifying statistically significant edges by comparison to, for example, a null-model [5] or edge weight distributions [6]. Although the aim of backbone extraction is reduction of data in order to simplify analysis, it is vital that the backbone still contains important and relevant information about the network. Hence, it is crucial that topological structures are preserved in the backbone [7].

Comparable approaches have been suggested for one-mode projections [8, 9]. A one-mode projection is the projection of a bipartite network onto a one-mode network. Bipartite networks consist of two disjoint sets of nodes U and V, where edges only connect nodes belonging to different sets. A bipartite network is usually represented by its biadjacency matrix B, where the element $b_{ij} = 1$ if node $u_i \in U$ is connected to node $v_j \in V$ and 0 otherwise. The one-mode projection of this bipartite network is constructed by dropping one node set and connecting two nodes of the remaining set if they share at least one neighbour in the bipartite network. Edges connecting the two nodes in the one-mode projection are associated with a weight, namely the number of previously shared neighbours.

Bipartite networks are far less studied than ordinary one-mode networks and hence many network measures are not applicable to bipartite structures. Although some measures have been redefined to suit the analysis of bipartite networks [10, 11, 12], the study of their projections is preferred in most cases [13].

Being able to infer the backbone of one-mode projections is especially important as the process of projection results inevitably in a dense and noisy one-mode network [8]. The extracted backbone could also be translated into a binary network, as it is often preferable to work with an unweighted version of the network [14].

Neal [9] introduced a stochastic degree sequence model to determine the significance of edges in weighted projections. In contrast to previous approaches, [9] takes the degrees of primary and secondary nodes in the original bipartite network into consideration. Determining the significant edges requires the generation of a reasonably large set of random bipartite networks and their one-mode projections. The random projections serve to determine the expected weight distribution of every edge in the network. Comparing the observed edge weights of the projections to their corresponding expected weight distribution allows the identification of significant edges.

The one-mode projection of a bipartite network is obtained by multiplying its biadjacency matrix B with its transpose. Hence, the process of projecting a bipartite net-

work is computationally expensive, running in $\mathcal{O}(|U|^2|V|)$ time and as [9] opines, a method to directly calculate edge weight distributions would be highly beneficial. In this paper we do exactly this, showing that the edge weights in most random one-mode projections are distributed according to a Poisson binomial distribution. We develop a fast method of extracting the backbone of a one-mode projection that does not rely on the generation of random networks and their projections and hence significantly reduces computation time. We corroborate the accuracy of our method by comparing the weights thus obtained to the real weight distribution of the projections of several types of randomly generated bipartite networks.

Further, we demonstrate, by means of two real world networks, that backbone extraction aids in the detection of communities. A community is loosely defined to be a subgraph of a network with a relatively higher number of inner connections compared to the number of edges linking to nodes outside the subgraph. We show that communities are more pronounced in the backbone than in the projection. This is a very interesting observation, since many real world networks exhibit community structure.

2. The weight distribution of one-mode projections

In the following, we show that the weight distribution of most projected random bipartite networks follows a Poisson binomial distribution. We give some remarks on exceptions at the end of this section.

2.1. The Poisson Binomial Distribution

We begin with the necessary definitions and notation.

Definition 1. *[15] A Bernoulli trial is a random variable X with two possible outcomes, success or failure, and is associated with a success probability p.*

The probability of obtaining n successes in N independent Bernoulli trials, where each trial has success probability p, is given by the binomial distribution [15]:

$$P(X_1 + \cdots + X_N = n) = \left(\begin{array}{c} N \\ n \end{array} \right) p^n (1-p)^{N-n}. \quad (1)$$

If the N trials have varying probabilities p_i, where $i = 1, \ldots, N$, the sum of the independent, non-identically distributed random variables $X_1, \ldots, X_N$ is given by the Poisson binomial distribution [16].

Let $\mathcal{S}_n$ be the set of all combinations of n distinct integers chosen from $\{1, \ldots, N\}$ and let $S_1, \ldots, S_{|\mathcal{S}_n|}$ be the elements of $\mathcal{S}_n$. Let s denote an element of the subset S_j, where $1 \leq j \leq |\mathcal{S}_n|$ and let $\bar{S}_j$ denote the complement of S_j with respect to $\{1, \ldots, N\}$. Then the probability density function of the Poisson binomial random variable $Z_X = \sum_{i=1}^{N} X_i$ is given by

$$P(Z_X = n) = \sum_{j=1}^{|\mathcal{S}_n|} \prod_{s \in S_j} p_s \prod_{\bar{s} \in \bar{S}_j} (1 - p_{\bar{s}}). \quad (2)$$

2.2. Approximation of the weight distribution

We now look at the use of the Poisson binomial distribution to approximate the weight distribution.

Definition 2. *[17, Definition 1.4] Let $a_0, a_1, a_2, \ldots$ be an arbitrary (infinite) sequence of numbers. The generating function for this sequence is the expression $\sum_{n=0}^{\infty} a_n x^n$.*

Let $\mathcal{B}$ be a bipartite network with the two disjoint node sets U and V, where U is called the primary set and V the secondary set. Let p_j denote the probability that a node $u \in U$ has degree j and q_k denote the probability that a node $v \in V$ has degree k.

By Definition 2, $f(x) = \sum_{j=0}^{\infty} p_j x^j$ is the probability generating function of the primary node degrees and $g(x) = \sum_{k=0}^{\infty} q_k x^k$ is the probability generating function of the secondary node degrees. It is important to note that

$$f(1) = \sum_{j=0}^{\infty} p_j = 1 \quad (3)$$

and

$$\left(x \frac{d}{dx} \right)^n f(x) \Big|_{x=1} = \sum_{j=0}^{\infty} j^n p_j = \langle j^n \rangle, \quad (4)$$

where $\langle j \rangle$ denotes the average node degree of the primary node set [18, Chapter 13]. Note that $\langle j^n \rangle \neq \langle j \rangle^n$.

The probability of an edge connecting a primary node to a secondary node is then determined by dividing the product of their degrees by the number of edges in the network. The number of edges m in a bipartite network is given by

$$m = |U|\langle j \rangle = |V|\langle k \rangle. \quad (5)$$

If $\pi_{uu'v}$ denotes the probability that two primary nodes u and u' are connected to a secondary node v, then given that $deg(u) = j_u$, $deg(u') = j_{u'}$ and $deg(v) = k_v$,

$$\begin{aligned} \pi_{uu'v} &= \frac{j_u j_{u'} k_v (k_v - 1)}{m(m-1)} \\ &= \frac{j_u j_{u'} k_v (k_v - 1)}{|U|^2 \langle j \rangle^2 - |U| \langle j \rangle}. \end{aligned} \quad (6)$$

Averaging j_u and $j_{u'}$ over the probabilities of an edge attaching to any node of degree j_u and $j_{u'}$ respectively, results in the probability π_v that any two primary nodes are connected to a particular secondary node v of degree k_v. Note that p_j is the fraction of nodes with degree j in the primary node set. Thus, the number of primary nodes with degree j is equal to $|U|p_j$ and the probability that an

edge originating at node v connects to a primary node of degree j is $|U|jp_j/m = jp_j/\langle j\rangle$ [18, Chapter 13]. Hence,

$$\begin{aligned}\pi_v &= \sum_{j_u,j_{u'}} \frac{j_u^2 p_{j_u} j_{u'}^2 p_{j_{u'}} k_v(k_v-1)}{(|U|^2\langle j\rangle^2 - |U|\langle j\rangle)\langle j\rangle^2} \\ &= \frac{k_v(k_v-1)}{|U|^2\langle j\rangle^4 - |U|\langle j\rangle^3}\left[\sum_{j=0}^{\infty} j^2 p_j\right]^2 \\ &= \frac{\langle j^2\rangle^2 k_v(k_v-1)}{|U|^2\langle j\rangle^4 - |U|\langle j\rangle^3}. \end{aligned} \quad (7)$$

π_v is the probability of the Bernoulli random variable X_v indicating the existence of a connection between two primary nodes via a particular secondary node v.

It follows that the probability of a randomly chosen edge in the one-mode projection having weight ω is given by $P\left(\Omega_X = \sum_{v=1}^{|V|} X_v = \omega\right) = \sum_{j=1}^{|\mathcal{S}_\omega|} \prod_{s\in S_j} \pi_s \prod_{\bar{s}\in\bar{S}_j}(1-\pi_{\bar{s}})$, where $\mathcal{S}_\omega$ is the set of all combinations of ω integers chosen from $\{1,\ldots,|V|\}$.

Since $P(\Omega_X = \omega)$ is hard to compute, we use the Poisson approximation instead:

$$P(\Omega_X = \omega) \approx \frac{\mu^\omega e^{-\mu}}{\omega!}, \quad (8)$$

where $\mu = \sum_{v=1}^{|V|} \pi_v$.

The error of the Poisson approximation to $P(\Omega_X = \omega)$ is $\epsilon_\omega < 2\sum_{i=1}^{|V|} \pi_v^2$ and is small if the number of expected successes is small [19]. Since most real world networks are sparse, the Poisson approximation estimates the probability of weight ω very well.

Notice that μ is also easily calculated:

$$\begin{aligned}\mu &= \sum_{v=1}^{|V|} \pi_v \\ &= \frac{\langle j^2\rangle^2}{|U|^2\langle j\rangle^4 - |U|\langle j\rangle^3} \sum_{v=0}^{|V|} k_v(k_v-1) \\ &= \frac{|V|\langle j^2\rangle^2\left(\langle k^2\rangle - \langle k\rangle\right)}{|U|^2\langle j\rangle^4 - |U|\langle j\rangle^3}. \end{aligned} \quad (9)$$

2.3. Determining probabilities of individual connections

When extracting the backbone of a network, one is interested in the probability of observing a connection with a certain weight between two nodes u and u' in the projection, denoted by $P_{uu'}(\Omega_X = \omega)$.

If $\pi_{uu'v}$ is small for every $v = 1,\ldots,|V|$, the Poisson approximation may be used, with

$$\begin{aligned}\mu &= \sum_{v=1}^{|V|} \pi_{uu'v} \\ &= \sum_{v=1}^{|V|} \frac{j_u j_{u'} k_v(k_v-1)}{|U|^2\langle j\rangle^2 - |U|\langle j\rangle} \\ &= \frac{|V| j_u j_{u'}\left(\langle k^2\rangle - \langle k\rangle\right)}{|U|^2\langle j\rangle^2 - |U|\langle j\rangle}. \end{aligned} \quad (10)$$

In bipartite networks where some nodes have a very high degree, it is often found that the probability of a connection between two individual nodes is very high, resulting in large approximation errors. In such situations, instead of the Poisson approximation, we use the normal approximation:

$$P_{uu'}(\Omega = \omega) \approx \frac{1}{\sigma\sqrt{2\pi}} e^{-(\omega-\mu)^2/(2\sigma^2)}, \quad (11)$$

where μ is given by Eq. (10) and

$$\begin{aligned}\sigma &= \left[\sum_{v=1}^{|V|} \pi_{uu'v}(1-\pi_{uu'v})\right]^{1/2} \\ &= \left[\sum_{v=1}^{|V|} \pi_{uu'v} - \sum_{v=1}^{|V|} \pi_{uu'v}^2\right]^{1/2} \\ &= \left[\mu - \sum_{v=1}^{|V|}\left(\frac{j_u j_{u'} k_v(k_v-1)}{|U|^2\langle j\rangle^2 - |U|\langle j\rangle}\right)^2\right]^{1/2} \\ &= \left[\mu - \frac{|V| j_u^2 j_{u'}^2\left(\langle k^4\rangle - 2\langle k^3\rangle + \langle k^2\rangle\right)}{|U|^4\langle j\rangle^4 - 2|U|^3\langle j\rangle^3 + |U|^2\langle j\rangle^2}\right]^{1/2} \end{aligned} \quad (12)$$

In order to demonstrate the accuracy of our approximation, we considered bipartite networks from all 25 possible permutations of the following degree distributions: Delta function, uniform, normal, exponential and power law distributions. For each permutation, we generated 100 random bipartite networks of the same size and projected

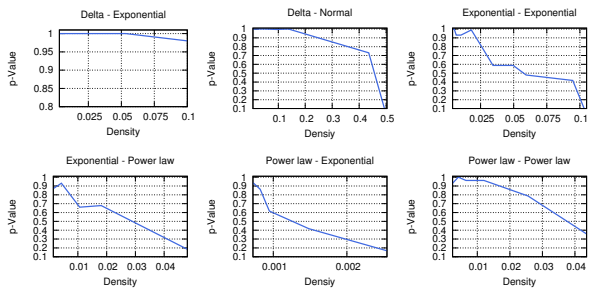

Figure 1: Results of the KS tests for six of the 25 tested cases.

each onto a one-mode network to determine their average weight distribution. We used the Kolmogorov-Smirnov (KS) test to compare the observed average weight distribution with the approximation. In addition, we tested our approximation for robustness, by changing the distribution parameters. The results of the KS tests suggest that the approximation estimates the expected weight distribution of a projected bipartite network extremely well and is robust against variation of the distribution parameters. As expected, our approximation performs poorly only for very dense networks, with p-values falling below 0.5. As most real world networks are extremely sparse such cases are rarely observed [20, 18]. Figure 1 shows the results of the KS tests for six commonly observed pairs of degree distributions.

In the extraction of the backbone of a projected bipartite network, we are interested in the probabilities of connections between individual nodes rather than the complete weight distribution. The weight probability distribution of an edge between two individual nodes u and u' is given by Eq. (11). Since $P_{uu'}(\Omega = \omega)$ has to be computed for every edge in the network, the computation time amounts to $\mathcal{O}(|U|^2)$. The greatest advantage of our method, however, is that it is not necessary to generate any random networks, saving the time required to generate and then project hundreds or thousands of networks. Since every single projection runs in $\mathcal{O}(|U|^2|V|)$ time our approach greatly simplifies and speeds up the process of extracting the backbone of a one-mode projection.

3. Applications

Next, we demonstrate our approach by extracting the backbone of two projected networks, the 108$^{\text{th}}$ U.S. Senate network [21, 22] and the MovieLens Tag Genome network [23]. The first network has previously been analysed, revealing that the backbone of the projected bipartite network contains two communities of democrats and republicans respectively [9]. Our work compares the performance of community detection algorithms using the one-mode projection and the backbone. We show that a higher modularity is achieved if the backbone is used as an input.

There has been much interest in the detection of network communities and a large body of literature exists on algorithms detecting them. We choose to use the approach suggested in [24], using eigenvectors, implemented in the R programming language [25]. This eigenvector community detection algorithm aims to divide the input network into groups such that the modularity function Q, obtained by subtracting the number of expected edges within communities from the observed number of connections within communities, is maximised. Rewriting Q in terms of matrices allows one to view the optimisation problem as a spectral problem, reducing its complexity, as community structure is often encoded in the first few eigenvectors of the modularity matrix. For more details see [24].

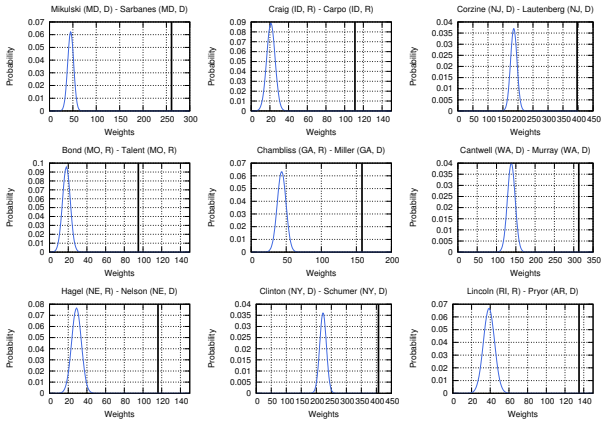
Mikulski (MD, D) - Sarbanes (MD, D)
Craig (ID, R) - Carpo (ID, R)
Corzine (NJ, D) - Lautenberg (NJ, D)
Bond (MO, R) - Talent (MO, R)
Chambliss (GA, R) - Miller (GA, D)
Cantwell (WA, D) - Murray (WA, D)
Hagel (NE, R) - Nelson (NE, D)
Clinton (NY, D) - Schumer (NY, D)
Lincoln (RI, R) - Pryor (AR, D)
Probability
Weights

Figure 2: The weight probability distributions of the nine most significant edges in the senator-senator projection, where the observed weight is greater than expected. The blue curves show the approximated probability distributions, the black vertical bars mark the observed weight in the weighted one-mode projection of the 108$^{\text{th}}$ U.S. Senate network.

3.1. 108th U.S. Senate Data

The U.S. Senate together with the House of Representatives builds the U.S. Congress. In every state two senators are voted into the senate and may introduce a piece of legislation, called bill. Bills can be co-sponsored by other members of the senate. The U.S. Senate data set may be represented as a bipartite network with 100 primary nodes, the senators, and 7,804 secondary nodes, the bills. An edge indicates that a senator has sponsored or co-sponsored a bill. This data set is publicly available and can be downloaded from http://jhfowler.ucsd.edu/cosponsorship.htm.

In order to extract the backbone of the senator-senator network, we need to find the weight probability distribution for every edge between all possible pairs of senators, 4,950 in total. If the observed edge weight between two senators is significantly larger than expected, the unweighted edge is included in the backbone. Here, an edge is included in the backbone if its weight is greater than the mean plus three standard deviations of the approximated distribution. The threshold of three standard deviations corresponds to a 95% confidence interval. This threshold is arbitrary and may be chosen differently. We found that lower thresholds attained lower modularities, whereas higher thresholds removed a larger number of edges, leading to the risk of deleting connections that are key to the network's topology. Although higher thresholds achieved higher modularities, the network simultaneously became more fragmented. For instance, choosing an extensive threshold of ten standard deviations resulted in a modularity of 0.57 and 27 groups of senators many of whom were isolated nodes. Albeit that the threshold is arbitrary, it has to be carefully chosen such that a high modularity is achieved while not fragmenting the network to the point of loss of important topological features.

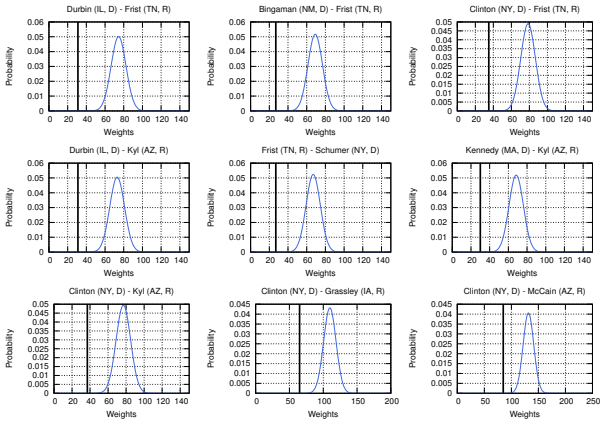
Durbin (IL, D) - Frist (TN, R)
Bingaman (NM, D) - Frist (TN, R)
Clinton (NY, D) - Frist (TN, R)
Durbin (IL, D) - Kyl (AZ, R)
Frist (TN, R) - Schumer (NY, D)
Kennedy (MA, D) - Kyl (AZ, R)
Clinton (NY, D) - Kyl (AZ, R)
Clinton (NY, D) - Grassley (IA, R)
Clinton (NY, D) - McCain (AZ, R)
Probability
Weights

Figure 3: The weight probability distributions of nine edges in the senator-senator projection, where the observed weight is smaller than expected (not included in the backbone). These edges represent political antagonisms. The blue curves show the approximated probability distributions, the black vertical bars mark the observed weight in the weighted one-mode projection of the 108$^{\text{th}}$ U.S. Senate network.

Plots of the weight distributions of some of the most significant edges are shown in Figs. 2 and 3.

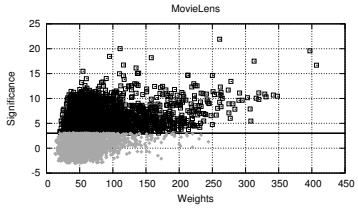

Figure 4: Edge significances in the U.S. Senate network plotted against their corresponding edge weights show no correlation between the two. Significant edges are represented by black squares.

Determining edge significance by comparing each edge weight individually to its expected distribution ensures that edges with high weights are not chosen over edges with low weights, for inclusion in the backbone [9]. Figure 4 illustrates that edge weights and edge significance are, with a correlation coefficient of 0.368, weakly correlated. Black squares indicate edges that were identified as being significant connections. Edge significance is calculated by subtracting the mean of the individual edge distribution from the observed weight and dividing by the distribution's standard deviation. It follows that the significance of edges cannot be trivially determined by removing all edges with a weight below a certain threshold. Consider for instance two senators with low degrees and assume they cosponsor the maximum possible number of bills, given by the minimum of the two degrees. Despite the edge connecting the two senators in the one-mode projection having a low weight, their relationship would be considered significant. Trivial edge removal may therefore lead to deletion of connections within communities and can never achieve high modularities, see Fig. 5.

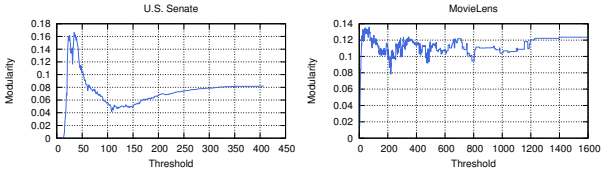

Figure 5: Extracting the backbone of the U.S Senate and the MovieLens networks by removing edges with weights under a certain threshold demonstrates that the increase in modularity achieved by our method (see below) is not trivial.

If the backbone is extracted by means of the method introduced in the previous section, besides redundant information being discarded, a more pronounced community structure is revealed. The majority of the edges connect senators from the same party with relatively few edges connecting senators from different parties (see Fig. 6, bottom left plot). Figure 6 shows the adjacency matrices of the binary projection (top left), the weighted projection (top right) and the backbone (bottom left) with a black square indicating the presence of an edge. The identified communities are highlighted by coloured rectangles.

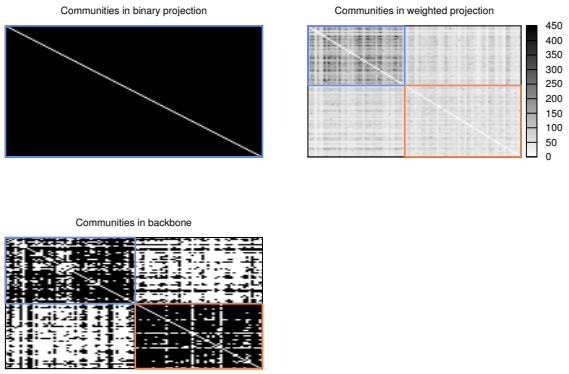

Figure 6: 108$^{\text{th}}$ U.S. Senate network: The adjacency matrices of the binary projection, the weighted projection and the backbone (see label of each). The identified communities are highlighted by coloured rectangles. As is evident, the backbone extraction preserves the intra community links, whereas many of the inter community connections are discarded. Senators are ordered according to their community association (first democrats then republicans) and then alphabetically by last name.

The binary one-mode projection (Fig. 6, top left) was the complete graph K_{100} and hence communities could not be identified. Consideration of the edge weights (Fig.

6, top right) resulted in the detection of two communities. Two communities were also detected in the backbone (Fig. 6, bottom left), however, they were more pronounced than in the weighted projection. The backbone achieved the highest modularity of 0.24 compared to 0.082 in the weighted projection. Note that the modularity ranges between 0 and 1. 94% of the senators associated with the first community were democratic senators, whereas 96% of the senators associated with the second community were republican. The republican senators who were associated with the first community were Lincoln Chafee, Susan M. Collins and Olympia J. Snowe, whereas the democratic senators associated with the second community were Kent Conrad and Zell Miller.

Researching these senators revealed the following: Lincoln Chafee was a member of the Republican Party until 2007, when he became an independent, later, in 2013, joining the Democratic Party. Susan Collins, a known moderate member of the Republican Party, is considered bipartisan. Like Collins, Olympia Snowe is also known to be strongly bipartisan. Kent Conrad was found to be more conservative than most other democratic politicians, hence his association with the second community of mostly republican party members. Zell Miller was also found to be conservative. In 2004 he backed President George W. Bush over the democratic nominee.

3.2. MovieLens Tags

The MovieLens Tag Genome dataset was collected by the University of Minnesota and is available for download at http://grouplens.org/datasets/movielens/tag-genome/. This dataset contains 9,734 movies and 1,128 tags. Tags are words assigned to movies by users of the MovieLens website. Users may use as a tag any word that they feel best describes a movie. Edges connect tags to movies and record the strength of the association of a particular movie with a particular tag. Edge weights range between 0 and 1, where 1 indicates strong relevance. In our analysis we considered a subset, including the 100 most popular tags (popularity as determined by members of the GroupLens research group, who also collected the data, http://grouplens.org/). Edges were only included if the tag relevance was greater or equal to 0.5, resulting in a network of 9,734 movies and 100 tags.

Extracting the backbone of the tag-tag projection should not only remove redundant information but additionally preserve connections between similar tags and remove those between unrelated tags, yielding more pronounced communities. For the same reasons as given for the U.S. Senate network, we chose a threshold of three standard deviations for the backbone extraction. Figure 7 shows the adjacency matrices of the binary projection (top left), the weighted projection (top right) and the backbone (bottom left). The community detection algorithm was able to identify two communities in the binary projection with a modularity of 0.018, three communities in the weighted projection with modularity 0.13 and five communities in the backbone with modularity 0.26.

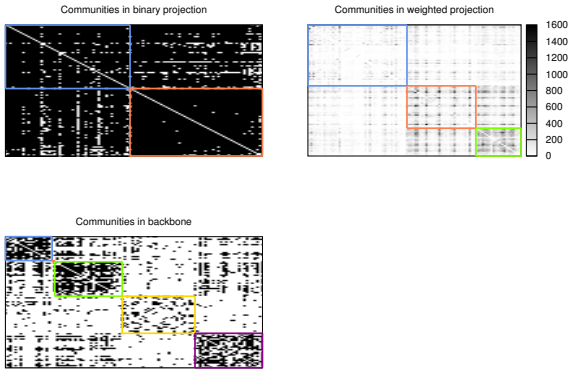

Figure 7: MovieLens Tag Genome network: The adjacency matrices of the binary projection, the weighted projection and the backbone (see label of each). The identified communities are highlighted by coloured rectangles. Again, the backbone extraction preserves the intra community connections, discarding many of the inter community connections. More information on the community membership of the different tags can be found in the supplementary movie TagNetwork.mp4.

Interestingly one of the tags is isolated in the backbone, forming a community by itself (see Fig. 8). This isolated tag, labelled boring, did not form any significant connections to other nodes in the network. The other four communities each contain tags that are very similar, for instance, the tags comedy, funny, animation, satire and pixar are members of the same community (see Fig. 9).

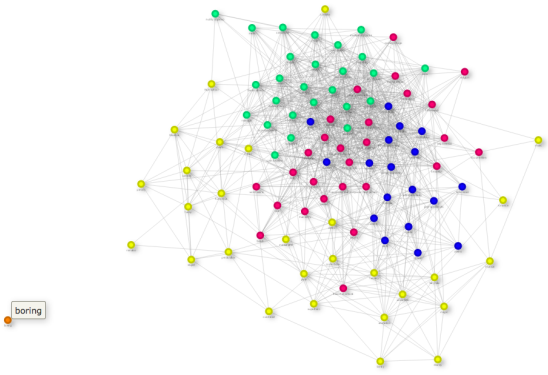

Figure 8: The backbone network of the Tag-tag projection shows an isolated node that forms a community by itself. The different colours represent community membership.

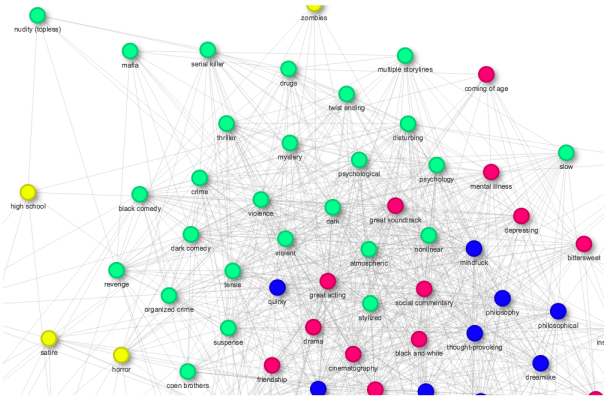

Figure 9: Looking closer at one of the communities in the backbone, we find that similar tags tend to form significant connections.

4. Conclusion

Since large complex networks are often difficult to analyse, it is important to develop efficient methods that can identify redundant information to shrink the data volume. In this paper, we have shown a simple and effective way of extracting the backbone of a one-mode projection, making use of the fact that the edge weights in most projected random networks follow a Poisson binomial distribution, thus drastically reducing the computation time needed for backbone extraction.

We used two different real world networks to demonstrate the backbone as an aid to detection of communities. In the process of extracting the backbone, connections between nodes of the same community are preserved, whereas inter community links are mostly insignificant and hence discarded. Thus the backbone displays more pronounced communities leading to detection algorithms achieving higher modularity scores.

References

[1] J. B. Glattfelder, S. Battiston, Backbone of complex networks of corporations: The flow of control, Physical Review E 80 (2009) 1–12.

[2] P. B. Slater, A two-stage algorithm for extracting the multiscale backbone of complex weighted networks, Proceedings of the National Academy of Sciences 106 (2009) E66.

[3] M. Scutari, R. Nagarajan, Identifying significant edges in graphical models of molecular networks, Artificial Intelligence in Medicine 57 (2013) 207–217.

[4] X. Zhang, Z. Zhang, H. Zhao, Q. Wang, J. Zhu, Extracting the globally and locally adaptive backbone of complex networks., PLoS One 9 (2014) e100428.

[5] M. A. Serrano, M. Boguñá, A. Vespignani, Extracting the multiscale backbone of complex weighted networks., Proceedings of the National Academy of Sciences of the United States of America 106 (2009) 6483–6488.

[6] N. J. Foti, J. M. Hughes, D. N. Rockmore, Nonparametric sparsification of complex multiscale networks, PLoS One 6 (2011).

[7] Q.-M. Zhang, A. Zeng, M.-S. Shang, Extracting the information backbone in online system., PLoS One 8 (2013) e62624.

[8] Z. Neal, Identifying statistically significant edges in one-mode projections, Social Network Analysis and Mining 3 (2013) 915–924.

[9] Z. Neal, The backbone of bipartite projections: Inferring relationships from co-authorship, co-sponsorship, co-attendance and other co-behaviors, Social Networks 39 (2014) 84–97.

[10] S. P. Borgatti, M. G. Everett, Network analysis of 2-mode data, Social Networks 19 (1997) 243–269.

[11] J. Liebig, A. Rao, Identifying influential nodes in bipartite networks using the clustering coefficient, in: Proceedings of the 2014 Tenth International Conference on Signal-Image Technology and Internet-Based Systems, 2014, pp. 323–330.

[12] J. Liebig, A. Rao, Predicting item popularity: Analysing local clustering behaviour of users, Physica A: Statistical Mechanics and its Applications 442 (2016) 523–531.

[13] M. E. J. Newman, The structure of scientific collaboration networks., Proceedings of the National Academy of Sciences of the United States of America 98 (2001) 404–409.

[14] M. Latapy, C. Magnien, N. D. Vecchio, Basic notions for the analysis of large two-mode networks, Social Networks 30 (2008) 31–48.

[15] C. Forbes, M. Evans, N. Hastings, B. Peacock, Statistical distributions, John Wiley & Sons, Inc., Hoboken, 2011.

[16] Y. H. Wang, On the numer of successes in independent trials, Statistica Sinica 2 (1993) 295–312.

[17] S. K. Lando, Lectures on generating functions, American Mathematical Society, Providence, 2003.

[18] M. E. J. Newman, Networks: An introduction, 1 ed., Oxford University Press, Oxford, 2010.

[19] L. Le Cam, An approximation theorem for the Poisson binomial distribution, Pacific Journal of Mathematics 10 (1960) 1181–1197.

[20] KONECT, Networks, 2014. URL: `http://konect.uni-koblenz.de/networks/`.

[21] J. H. Fowler, Connecting the congress: A study of cosponsorship networks, Political Analysis 14 (2006) 456–487.

[22] J. H. Fowler, Legislative cosponsorship networks in the US House and Senate, Social Networks 28 (2006) 454–465.

[23] J. Vig, S. Sen, J. Riedl, The tag genome: Encoding community knowledge to support novel interaction, ACM Transactions on Interactive Intelligent Systems (TiiS) 2 (2012) 13.

[24] M. E. J. Newman, Finding community structure in networks using the eigenvectors of matrices, Physical Review E 74 (2006) 036104.

[25] G. Csárdi, T. Nepusz, The igraph software package for complex network research., InterJournal, Complex Systems 1695 (2006) 1–9.